Deep Transfer Learning for Single-Channel Automatic Sleep Staging with Channel Mismatch

Huy Phan[†], Oliver Y. Chén[*], Philipp Koch[‡], Alfred Mertins[‡], and Maarten De Vos[*]

[†]*School of Computing, University of Kent, United Kingdom*
[*]*Institute of Biomedical Engineering, University of Oxford, United Kingdom*
[‡]*Institute for Signal Processing, University of Lübeck, Germany*

Abstract—Many sleep studies suffer from the problem of insufficient data to fully utilize deep neural networks as different labs use different recordings set ups, leading to the need of training automated algorithms on rather small databases, whereas large annotated databases are around but cannot be directly included into these studies for data compensation due to channel mismatch. This work presents a deep transfer learning approach to overcome the channel mismatch problem and transfer knowledge from a large dataset to a small cohort to study automatic sleep staging with single-channel input. We employ the state-of-the-art *SeqSleepNet* and train the network in the source domain, *i.e.* the large dataset. Afterwards, the pretrained network is finetuned in the target domain, *i.e.* the small cohort, to complete knowledge transfer. We study two transfer learning scenarios with slight and heavy channel mismatch between the source and target domains. We also investigate whether, and if so, how finetuning entirely or partially the pretrained network would affect the performance of sleep staging on the target domain. Using the Montreal Archive of Sleep Studies (MASS) database consisting of 200 subjects as the source domain and the Sleep-EDF Expanded database consisting of 20 subjects as the target domain in this study, our experimental results show significant performance improvement on sleep staging achieved with the proposed deep transfer learning approach. Furthermore, these results also reveal the essential of finetuning the feature-learning parts of the pretrained network to be able to bypass the channel mismatch problem.

Index Terms—Automatic sleep staging, deep learning, transfer learning, SeqSleepNet

I. INTRODUCTION

Deep learning has been successfully applied to numerous domains and has received much attention from the sleep research community. With its ability in automatic representation learning to make sense of large datasets, deep learning has improved the performance of automatic sleep staging significantly, reaching an accuracy rate on par with manual scoring by sleep experts [1], [2]. However, this expert-level performance is only obtainable when the studied cohort is of large size, *i.e.* hundreds to thousands of subjects [1], [2]. The reason for this is that deep neural networks usually require a large amount of data to train. In practice, many sleep studies possess a small cohort, such as a few dozens of subjects [3], [4], [5] sometimes even less [6], [7], particularly those studies related to sleep disorders [7], [8] or exploring new channel positions [6]. As a consequence, the small amount

Corresponding author: `h.phan@kent.ac.uk`.

of data makes deep neural networks underperform in these sleep studies. One possibility to remedy the lack of data is to make use of the large databases which are manually annotated and publicly available. However, these databases cannot be simply added into these studies due to channel mismatch. The problem of channel mismatch arises when different studies uses different channel layouts [2] or when novel electrode placements might be explored [6]. Moreover, it also happens when a study investigates a particular sleep abnormality, poor performance can be obtained when the automated diagnostic procedure is only trained on healthy volunteers [7].

We propose a transfer learning approach to circumvent the channel-mismatch problem, enabling efficient knowledge transfer from a large database to study single-channel sleep staging in a much smaller dataset with deep neural networks. We use Montreal Archive of Sleep Studies (MASS) database with 200 subjects [9] as the source domain and the Sleep-EDF Expanded database with 20 subjects [3], [4] as the target domain in this study. *SeqSleepNet*, a *sequence-to-sequence* model which was recently shown to achieve state-of-the-art results on sleep staging in our recent work [1], is employed as the base model. The network is firstly trained with the source domain data and is then finetuned with the target domain data. We study two scenarios: (1) same-modality transfer learning with slight channel mismatch when the source domain is an EEG channel and the target domain is another EEG channel, (2) cross-modality transfer learning with heavy channel mismatch when the source domain is an EEG channel and the target domain is an EOG channel. Furthermore, we investigate different finetuning strategies to gain insight into deep transfer learning for sleep staging under different channel mismatch conditions.

II. SLEEP STAGING DATABASES

A. Montreal Archive of Sleep Studies (MASS)

We employed the public MASS dataset [9] as the source domain in this study. This dataset consists of whole-night recordings from 200 subjects aged between 18 and 76 years (97 males and 103 females). Manual annotation was done on each epoch of the recordings by sleep experts according to the AASM standard [10] (SS1 and SS3 subsets) or the R&K standard [11] (SS2, SS4, and SS5 subsets). As in [12], [1], we converted different annotations into five sleep stages {W, N1,

N2, N3, and REM} and expanded 20-second epochs into 30-second ones by including 5-second segments before and after each epoch. We adopted the C4-A1 EEG channel (C4-A1) as the source domain data. The signals, originally sampled at 256 Hz, were downsampled to 100 Hz.

B. Sleep-EDF Expanded (Sleep-EDF)

We adopted the Sleep Cassette (SC) subset of the Sleep-EDF dataset [3], [4] as the target domain in this study. It consists of 20 subjects aged 25-34. The PSG recordings, sampled at 100 Hz, of two subsequent day-night periods were available for each subject, except for one subject (subject 13) who had only one-night data. Each 30-second epoch of the recordings was manually labelled by sleep experts according to the R&K standard [11] into one of eight categories {W, N1, N2, N3, N4, REM, MOVEMENT, UNKNOWN}. Similar to previous works [13], [14], [15], N3 and N4 stages were merged into a single stage N3. MOVEMENT and UNKNOWN categories were excluded. We explored the Fpz-Cz EEG channel and the EOG (horizontal) channel in our same-modality and cross-modality transfer learning experiments. Only the in-bed parts (from *lights off* time to *lights on* time) of the recordings were included as recommended in [16], [13], [14], [12].

III. SEQSLEEPNET-BASED TRANSFER LEARNING FOR SINGLE-CHANNEL AUTOMATIC SLEEP STAGING

A. Single-channel SeqSleepNet

SeqSleepNet [1] has demonstrated state-of-the-art performance on the MASS dataset with multi-channel input. We adapt it to study single-channel sleep staging here. The network was proposed to deal with sequence-to-sequence sleep staging, *i.e.* mapping a sequence of multiple consecutive epochs to a sequence of sleep stage labels at once [1]. Formally, sequence-to-sequence sleep staging is formulated to maximize the conditional probability $p(\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_L \,|\, \mathbf{S}_1, \mathbf{S}_2, \ldots, \mathbf{S}_L)$ where $\mathbf{S}_1, \mathbf{S}_2, \ldots, \mathbf{S}_L$ denote L consecutive epochs in the input sequence and $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_L$ represent L one-hot encoding vectors of the ground-truth output labels corresponding to the input epochs.

The overall architecture of *SeqSleepNet* is illustrated in Fig 1(a). The network is composed of three main components: an attentional RNN (ARNN) subnetwork for epoch-wise feature learning, a sequence-level RNN (SeqRNN) for long-term sequential modelling, and a softmax layer for sequence classification.

ARNN. The ARNN subnetwork, illustrated in Fig. 1(b), is shared between all epochs in the input sequence. It is the combination of a *filterbank* layer [17], an epoch-level bidirectional RNN (eRNN for short to distinguish it from the SeqRNN), and an attention layer [18]. The subnetwork receives a time-frequency image $\mathbf{S} \in \mathbb{R}^{F \times T}$ obtained from an epoch as input (i.e. an epoch $\mathbf{S}_i$ in the input sequence, cf. Fig. 1), where F and T denote the frequency and time dimensions, respectively. The input is firstly preprocessed by the filterbank layer of M filters ($M < F$) for frequency-dimension reduction and smoothing [17]. Afterwards, the
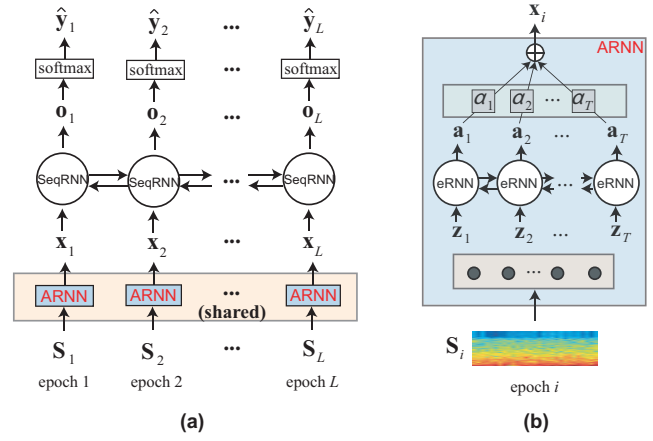

Fig. 1. (a) *SeqSleepNet*'s overall architecture and (b) its attentional RNN (ARNN) subnetwork.

resulting image is treated as a sequence of T local feature vectors (*i.e.* image columns) $(\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_T)$, which is then encoded into a sequence of output vectors $(\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_T)$, where

$$\mathbf{a}_t = \mathbf{W}_{ha}[\mathbf{h}_t^{\mathrm{b}} \oplus \mathbf{h}_t^{\mathrm{f}}] + \mathbf{b}_a, \quad 1 \leq t \leq T. \tag{1}$$

In (1), $\mathbf{W}_{ha}$ denotes the weight matrix, $\mathbf{b}_a$ denotes the bias, and $\oplus$ denotes vector concatenation. $\mathbf{h}_i^{\mathrm{b}}$ and $\mathbf{h}_i^{\mathrm{f}}$ represent backward and forward hidden state vectors of the backward and forward recurrent layers at time index t, respectively, and are computed as

$$\mathbf{h}_t^{\mathrm{f}} = \mathcal{H}(\mathbf{x}_t, \mathbf{h}_{t-1}^{\mathrm{f}}), \tag{2}$$

$$\mathbf{h}_t^{\mathrm{b}} = \mathcal{H}(\mathbf{x}_t, \mathbf{h}_{t+1}^{\mathrm{b}}), \ 1 \leq t \leq T. \tag{3}$$

Here, $\mathcal{H}$ denotes the hidden layer function which is realized by a Gated Recurrent Unit (GRU) cell [19]. In turn, the attention layer [18] learns to weight the output vectors and collectively form the feature vector $\mathbf{x}$:

$$\mathbf{x} = \sum_{t=1}^{T} \alpha_t \mathbf{a}_t. \tag{4}$$

Here, α_t is the attention weight at the time index t learned by the attention layer.

SeqRNN. Via the shared ARNN, the input sequence $(\mathbf{S}_1, \mathbf{S}_2, \ldots, \mathbf{S}_L)$ has been transformed into a sequence of epoch-wise feature vectors $(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_L)$. SeqRNN, which is essentially also a GRU-based bidirectional RNN, reads the induced sequence of feature vectors forward and backward to encode their sequential interactions into a sequence of output vectors $(\mathbf{o}_1, \mathbf{o}_2, \ldots, \mathbf{o}_L)$, where

$$\mathbf{o}_i = \mathbf{W}_{ho}[\tilde{\mathbf{h}}_i^{\mathrm{b}} \oplus \tilde{\mathbf{h}}_i^{\mathrm{f}}] + \mathbf{b}_o, \tag{5}$$

$$\tilde{\mathbf{h}}_i^{\mathrm{f}} = \mathcal{H}(\mathbf{o}_i, \tilde{\mathbf{h}}_{i-1}^{\mathrm{f}}), \tag{6}$$

$$\tilde{\mathbf{h}}_i^{\mathrm{b}} = \mathcal{H}(\mathbf{x}_i, \tilde{\mathbf{h}}_{i+1}^{\mathrm{b}}), \ 1 \leq i \leq L. \tag{7}$$

Softmax. The softmax layer eventually classifies the sequence of output vectors $(\mathbf{o}_1, \mathbf{o}_2, \ldots, \mathbf{o}_L)$ to produce the output sequence of sleep stage probabilities $(\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2, \ldots, \hat{\mathbf{y}}_L)$.

The network is trained to minimize the sequence classification loss [1] over N training sequences in the training data:

$$E(\boldsymbol{\theta}) = -\frac{1}{L}\sum_{n=1}^{N}\sum_{i=1}^{L}\mathbf{y}_i\log\left(\hat{\mathbf{y}}_i\left(\boldsymbol{\theta}\right)\right) + \frac{\lambda}{2}\|\boldsymbol{\theta}\|_2^2. \qquad (8)$$

Although being hierarchical, the network is trained in an end-to-end manner similar to [1].

IV. SEQSLEEPNET-BASED TRANSFER LEARNING

Insufficient training data is a common problem in many sleep studies as it is expensive and difficult to collect a large-scale and high-quality annotated sleep dataset. Meanwhile, there exist large sleep databases which are publicly available, consisting data from hundreds to thousands of subjects, such as MASS dataset [9]. Unfortunately, these datasets cannot be simply included into the private studies due to channel mismatch (see Section I). Transfer learning [20] relaxes the hypothesis that the training data must be independent and identically distributed (i.i.d.) from the test data, and therefore, holds promise to leverage these large amount of available data to tackle the problem of insufficient training data in the studies with a small number of subjects. We investigate deep transfer learning with *SeqSleepNet* as the base model, the MASS database as the source domain, and the Sleep-EDF database as the target domain. A *SeqSleepNet* will be used as a device to transfer the knowledge from the MASS database for sleep staging with the Sleep-EDF database as illustrated in Fig. 2.

We study the two transfer learning scenarios for single-channel sleep staging, taking into account different degrees of channel mismatch:

EEG↦EEG: In this same-modality transfer learning, *SeqSleepNet* is firstly trained on the C4-A1 EEG channel of the MASS database and is then transferred to the Fpz-Cz EEG channel of the Sleep-EDF database. This study is to showcase that even the source domain and the target domain are of the same modality, transfer learning is still necessary to overcome the small channel mismatch and to improve the sleep staging performance. Note that we adopt two different EEG channels for the source and target domains on purpose not only because the C4-A1 EEG channel is not available in the Sleep-EDF database but also because we want to enforce more channel mismatch here.

EEG↦EOG: In this cross-modality transfer learning, *SeqSleepNet* is firstly trained on the C4-A1 EEG channel of the MASS database and is then transferred to the horizontal EOG channel of the Sleep-EDF database. On the one hand, this study is to demonstrate that transfer learning is essential to tackle heavy channel mismatch between the source and target domains to improve sleep staging performance on the target domain. On the other hand, as EOG signals contain rich information from multiple sources, including ocular activity, frontal EEG activity, and EMG from cranial and eye muscles, they are promising candidates for a single-modality sleep staging system. Despite their potential, EOG signals have
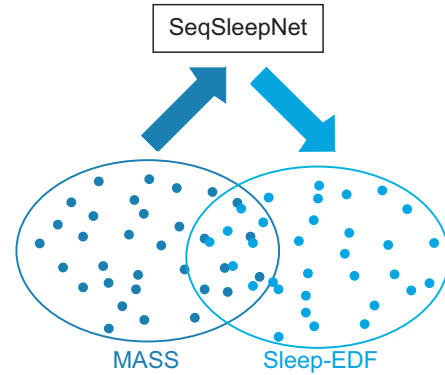

Fig. 2. *SeqSleepNet*-based transfer learning from the source domain (MASS) to the target domain (Sleep-EDF) with channel mismatch.

been mainly used as a complement for EEG signals in multi-modality sleep staging studies [1], [12]. Only a few studies have exploited standalone EOG signals for single-modality sleep staging [5], [21]. We therefore want to examine whether deep transfer learning could help sleep staging with standalone secondary EOG to achieve a performance level comparable to that using the primary EEG. Furthermore, due to the ease of electrode placements, they would be ideal for home-based sleep monitoring with wearable devices [6].

In order to accomplish transfer learning, the *SeqSleepNet* pretrained in the source domain is considered as a starting point in the target domain. We reuse the entire or parts of the pretrained network and finetune the rest with the target domain data to study the influence of finetuning different parts of the pretrained network to the sleep stage performance on the target domain. To that end, we examine directly transferring (*i.e.* the pretrained network is directly used in the target domain without finetuning), finetuning the entire pretrained network, and partially finetuning {softmax, softmax + ARNN, softmax + SeqRNN} combinations of the pretrained network while other layers stay fixed.

V. EXPERIMENTS

A. Experimental setup

While the entire 200 subjects of the MASS database (*i.e.* the source domain) was used for network pretraining, we conducted leave-one-subject-out cross validation on the Sleep-EDF database (*i.e.* the target domain) to access the efficiency of transfer learning on the target domain. At each iteration, 19 remaining subjects of Sleep-EDF were further divided into 15 subjects for finetuning the pretrained model and 4 subjects for validation. The performance over 20 cross-validation folds will be reported.

B. Time-frequency input

As *SeqSleepNet* receives time-frequency images as input, each 30-second epoch (*i.e.* EEG and EOG) was transformed into log-power spectra via short-time Fourier transform (STFT) with a window size of two seconds and 50% overlap, followed by logarithmic scaling. Hamming window and 256-point Fast Fourier Transform (FFT) were used. This results in

an image $\mathbf{S} \in \mathbb{R}^{F \times T}$ where $F = 129$ (the number of frequency bins), $T = 29$ (the number of spectral columns).

C. Network parameters

The implementation was based on *Tensorflow* framework [22]. For pretraining, the network was trained with the MASS dataset for 10 training epochs with a minibatch size of 32 sequences. For finetuning, a pretrained network was further trained with the Sleep-EDF training data for 10 training epochs. Early stopping was exercised, during which the finetuning was stopped if after 50 training steps there was not any improvement on the Sleep-EDF validation set. The network training and finetuning were performed using *Adam* optimizer with a learning rate of 10^{-4}.

The network was parametrized similarly to that in our previous work [1] (i.e. the number of filters M, the size of the GRUs' hidden state vectors, dropout, etc.) except that its input is single-channel here. We experimented with sequence length $L = 20$ epochs. Similar to [1], the sequences were sampled from the PSG recordings with a maximum overlap (i.e. $L - 1$ epochs) to generate all possible sequences from the training recordings for network training and finetuning. During testing, the test sequence was shifted by one to obtain an ensemble of classification decisions at each epoch of the test PSG recordings. Probabilistic multiplicative fusion was then carried out to aggregate these decisions to produce the final decision as in [1].

D. Experimental results

1) Results on **EEG**$\mapsto$**EEG** *transfer learning:* Table I shows the performance of the **EEG**$\mapsto$**EEG** transfer learning scenario. We also compare these results with those reported in the literature. The overall performance is reported using accuracy, macro F1-score (MF1), Cohen's kappa (κ), sensitivity, and specificity. Note that, as Sleep-EDF was used differently in previous works [12], [16], for compatible comparison, we only included those with a similar experimental setting (i.e. conducting independent testing with single FPz-Cz channel and in-bed data only [12]). A more comprehensive comparison can be found in [12].

The results in Table I show that, under the condition of small channel mismatch, direct transfer obtains a good accuracy level which is close to the best reported in previous work (i.e. Multitask 1-max CNN [12]). However, this accuracy is 1.0% lower than the *SeqSleepNet* trained from scratch with the target domain data only, highlighting the necessity of finetuning even with same-modality transfer learning. As expected, entire or partial finetuning the pretrained *SeqSleepNet* results in significant performance gain over the naive direct transfer strategy. Absolute gains of 0.4%, 4.0%, 4.2%, and 4.3% on overall accuracy are obtained when finetuning softmax, softmax + ARNN, softmax + SeqRNN, and all layers of the pretrained network, respectively. Moreover, the finetuning results significantly outperform the best result obtained in previous work, improving the single-channel sleep staging accuracy on the Sleep-EDF from 81.9% (i.e. Multitask 1-max CNN [12]) to 85.5% (i.e. finetuning entire network).

TABLE I
RESULTS OF THE **EEG**$\mapsto$**EEG** TRANSFER LEARNING IN COMPARISON WITH RESULTS REPORTED IN PREVIOUS WORK. FINETUNING IS ABBREVIATED AS FT FOR SHORT.

System	Transfer learning	Overall metrics				
		Acc.	κ	MF1	Sens.	Spec.
FT Entire network	Yes	**85.5**	**0.794**	**80.0**	**79.6**	**95.9**
FT Softmax + ARNN	Yes	85.4	0.792	79.9	79.6	95.9
FT Softmax + SeqRNN	Yes	85.2	0.789	79.6	78.7	95.8
FT Softmax	Yes	81.6	0.735	74.9	74.4	94.8
Direct transfer (no FT)	Yes	81.2	0.733	74.6	74.1	94.8
SeqSleepNet (scratch)	No	82.2	0.746	74.1	73.9	95.0
Multitask 1-max CNN [12]	No	81.9	0.740	73.8	73.9	95.0
1-max CNN [17]	No	79.8	0.720	72.0	—	—
Atentional RNN [23]	No	79.1	0.700	69.8	—	—
Deep auto-encoder [14]	No	78.9	—	73.3	—	—
Deep CNN [13]	No	74.8	—	69.8	—	—

TABLE II
RESULTS OF THE **EEG**$\mapsto$**EOG** TRANSFER LEARNING. FINETUNING IS ABBREVIATED AS FT FOR SHORT.

System	Transfer learning	Overall metrics				
		Acc.	κ	MF1	Sens.	Spec.
FT Entire Network	Yes	80.8	0.725	74.2	73.1	94.5
FT Softmax + ARNN	Yes	80.1	0.716	73.5	72.5	94.3
FT Softmax + SeqRNN	Yes	80.0	0.709	72.3	70.4	94.0
FT Softmax	Yes	73.2	0.605	62.5	60.7	91.9
Direct transfer (no FT)	Yes	51.1	0.300	42.5	42.9	85.9
SeqSleepNet (scratch)	No	78.5	0.688	68.3	67.4	93.6

The finetuning results also unveil that finetuning the softmax layer alone is not sufficient to overcome the channel-mismatch obstacle. Instead, it is important to additionally finetune those feature-learning layers, either the ARNN subnetwork for epoch-level feature learning or the SeqRNN for sequence-level feature learning or both collectively. In Fig. 3, we show the development curves of the test accuracy of four different cross-validation folds during finetuning. Earlier convergence and better generalization can be consistently seen when finetuning the feature-learning parts of the network in addition to the softmax layer.

2) Results on **EEG**$\mapsto$**EOG** *transfer learning:* Table II shows the results of the **EEG**$\mapsto$**EOG** transfer learning in comparison to that obtained by the *SeqSleepNet* trained from scratch using the target domain data only. Firstly, it can be seen that direct transfer does not work properly under the heavy channel-mismatch condition. It is understandable that differences in characteristics of the source EEG and the target EOG cause malfunction of the feature-learning parts of the pretrained *SeqSleepNet* in the target domain. As a consequence, finetuning is essential in this case. Secondly, the finetuning results exhibit similar patterns to the same-modality transfer learning case. That is, it is chief to finetune the feature-learning layers, either the ARNN subnetwork or the SeqRNN or both, alongside with the softmax layer. Finetuning leads to absolute gains of 22.1%, 28.9%, 29.0%, and 29.7% on overall accuracy compared to the naive direct transfer when softmax, softmax + ARNN, softmax + SeqRNN, and all layers of the pretrained network are finetuned, respectively.
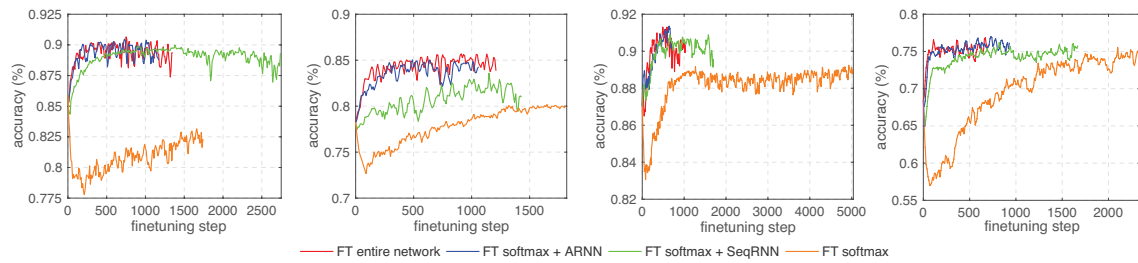
Fig. 3. Variation of the test accuracy as a function of the finetuning step.

More importantly, despite the heavy channel mismatch condition, transferring the knowledge of the source domain to the target domain brings up the accuracy by 2.3% over the network trained from scratch with the target domain data only. With the obtained accuracy of 80.8%, single-channel sleep staging with the secondary EOG reaches, and sometimes even outperforms, many deep-learning approaches based on the primary EEG presented in Table I. Therefore, this finding would allow one to explore the usage of EOG as an alternative for EEG in many applications, such as home-based sleep monitoring with wearable devices [6], [24].

VI. CONCLUSIONS

We propose a deep transfer learning approach using the state-of-the-art *SeqSleepNet* to borrow and transfer knowledge from the large MASS database (the source domain) to enable exploiting the modelling power of deep neural networks on single-channel sleep staging with the Sleep-EDF database (*i.e.* the target domain) with a small number of subjects. *SeqSleepNet* was firstly trained with the C4-A1 EEG channel of the MASS database and was then transferred to the Fpz-Cz EEG channel and the horizontal EOG channel of the Sleep-EDF database via finetuning. The experimental results showed that the sleep staging performance on the target domain was significantly improved in both same-modality and cross-modality transfer learning scenarios. In addition, we demonstrated that it is essential to finetune the feature-learning parts of the pretrained network to overcome the channel mismatch problem.

REFERENCES

[1] H. Phan, F. Andreotti, N. Cooray, O. Y. Chén, and M. De Vos, "SeqSleepNet: end-to-end hierarchical recurrent neural network for sequence-to-sequence automatic sleep staging," *IEEE Trans. Neural Systems and Rehabilitation Engineering (TNSRE)*, vol. 27, no. 3, pp. 400–410, 2019.

[2] J. B. Stephansen *et al.*, "Neural network analysis of sleep stages enables efficient diagnosis of narcolepsy," *Nature Communications*, vol. 9, no. 1, pp. 5229, 2018.

[3] B. Kemp, A. H. Zwinderman, B. Tuk, H. A. C. Kamphuisen, and J. J. L. Oberye, "Analysis of a sleep-dependent neuronal feedback loop: the slow-wave microcontinuity of the EEG," *IEEE Trans. on Biomedical Engineering*, vol. 47, no. 9, pp. 1185–1194.

[4] A. L. Goldberger *et al.*, "Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, pp. e215–e220, 2000.

[5] A. N. Olesen, J. A. E. Christensen, H. B. D. Sorensen, and P. J. Jennum, "A noise-assisted data analysis method for automatic eog-based sleepstage classification using ensemble learning," in *Proc. EMBC*, 2016, pp. 3769–3772.

[6] K. B. Mikkelsen *et al.*, "Machine-learning-derived sleep-wake staging from around-the-ear electroencephalogram outperforms manual scoring and actigraphy," *Journal of Sleep Research*, p. e12786, 2018.

[7] N. Cooray, F. Andreotti, C. Lo, M. Symmonds, M. T. M. Hu, and M. De Vos, "Detection of REM sleep behaviour disorder by automated polysomnography analysis," *arXiv preprint arXiv:1811.04662*, 2018.

[8] F. Andreotti, H. Phan, N. Cooray, C. Lo, M. T. M. Hu, and M. De Vos, "Multichannel sleep stage classification and transfer learning using convolutional neural networks," in *Proc. EMBC*, 2018, pp. 171–174.

[9] C. O'Reilly, N. Gosselin, J. Carrier, and T. Nielsen, "Montreal archive of sleep studies: An open-access resource for instrument benchmarking & exploratory research," *Journal of Sleep Research*, pp. 628–635, 2014.

[10] C. Iber, S. Ancoli-Israel, A. L. Chesson, and S. F. Quan, "The AASM manual for the scoring of sleep and associated events: Rules, terminology and technical specifications," *American Academy of Sleep Medicine*, 2007.

[11] J. A. Hobson, "A manual of standardized terminology, techniques and scoring system for sleep stages of human subjects," *Electroencephalography and Clinical Neurophysiology*, vol. 26, no. 6, pp. 644, 1969.

[12] H. Phan, F. Andreotti, N. Cooray, O. Y. Chén, and M. De Vos, "Joint classification and prediction CNN framework for automatic sleep stage classification," *IEEE Trans. Biomedical Engineering (TBME)*, vol. 66, no. 5, pp. 1285–1296, 2019.

[13] O. Tsinalis, P. M. Matthews, Y. Guo, and S. Zafeiriou, "Automatic sleep stage scoring with single-channel EEG using convolutional neural networks," *arXiv:1610.01683*, 2016.

[14] O. Tsinalis, P. M. Matthews, and Y. Guo, "Automatic sleep stage scoring using time-frequency analysis and stacked sparse autoencoders," *Annals of Biomedical Engineering*, vol. 44, no. 5, pp. 1587–1597, 2016.

[15] A. Supratak, H. Dong, C. Wu, and Y. Guo, "DeepSleepNet: A model for automatic sleep stage scoring based on raw single-channel EEG," *IEEE Trans. on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 11, pp. 1998–2008, 2017.

[16] S. A. Imtiaz and E. Rodriguez-Villegas, "An open-source toolbox for standardized use of PhysioNet Sleep EDF Expanded Database," in *Proc. EMBC*, 2015, pp. 6014–6017.

[17] H. Phan, F. Andreotti, N. Cooray, O. Y. Chén, and M. De Vos, "DNN filter bank improves 1-max pooling CNN for single-channel EEG automatic sleep stage classification," in *Proc. EMBC*, 2018, pp. 453–456.

[18] T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proc. EMNLP*, 2015, pp. 1412–1421.

[19] K. Cho, B. van Merrienboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *Proc. EMNLP*, 2014, pp. 1724–1734.

[20] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.

[21] S.-F. Liang *et al.*, "Development of an EOG-based automatic sleep-monitoring eye mask," *IEEE Trans. Instrum. Meas.*, vol. 64, no. 11, pp. 2977–2985, 2015.

[22] M. Abadi *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv:1603.04467*, 2016.

[23] H. Phan, F. Andreotti, N. Cooray, O. Y. Chén, and M. De Vos, "Automatic sleep stage classification using single-channel eeg: Learning sequential features with attention-based recurrent neural networks," in *Proc. EMBC*, 2018, pp. 1452–1455.

[24] D. Looney *et al.*, "Wearable in-ear encephalography sensor for monitoring sleep. preliminary observations from nap studies," *Annals of the American Thoracic Society*, vol. 13, no. 12, pp. 32–42, 2016.